

Nils-Eric Sahlin

SECONDARY LEVELS IN DECISION MAKING

Lund 1984



Digitized by the Internet Archive
in 2026

https://archive.org/details/IA41546724_0073

Organization LUND UNIVERSITY Department of Philosophy Kungshuset, Lundagård S-223 50 LUND, Sweden		Document name DOCTORAL DISSERTATION	
		Date of issue September 1984	
		CODEN: LUHFDA/(HFFT-1005)/1-28/(1984)	
Author(s) Nils-Eric Sahlin		Sponsoring organization	
Title and subtitle Secondary Levels in Decision Making			
Abstract This thesis is devoted to three topics. In section 1, a decision theory which allows the agent to account for the uncertainties of the probability assessments is developed and discussed. This is accomplished by modelling the beliefs about the states of nature by means of a class of probability measures. In order to represent the uncertainty of these beliefs a measure of epistemic reliability, a second order probability measure, is introduced. In section 2, a method for obtaining information about an agent's utility assessments is presented and compared with a classical method based on 50-50 game comparison. The method is based on the concept of second order preferences. Section 3, finally, contains a discussion of some fundamental problems of the utility theory measure of risk suggested by Arrow and Pratt. It is suggested that this measure is not an adequate measure of risk. A new measure based on the concept of level of aspiration and the idea that a utility theory measure of risk is best defined as a comparison between two utility functions, one 'risk-free' and the other 'affected by risk', is presented and discussed.			
Key words probability; second order probability; preference; second order preferences; decision theory; utility theory; risk; epistemic risk; level of aspiration			
Classification system and/or index terms (if any)			
Supplementary bibliographical information			Language English
ISSN and key title			ISBN
Recipient's notes		Number of pages 28+103	Price
		Security classification	

Distribution by (name and address) Nils-Eric Sahlin, Department of Philosophy, Lund University, Kungshuset, Lundagård, S-223 50 Lund, Sweden

I, the undersigned, being the copyright owner of the abstract of the above-mentioned dissertation, hereby grant to all reference sources permission to publish and disseminate the abstract of the above-mentioned dissertation.

Signature Nils-Eric Sahlin

Date September 4, 1984

SECONDARY LEVELS IN DECISION MAKING

AV

NILS-ERIC SAHLIN

Fil. kand., Lu

AKADEMISK AVHANDLING

som för avläggande av filosofie doktorsexamen
vid Humanistiska fakulteten vid universitetet i Lund
kommer att offentligen försvaras å Carolinasalen
tisdagen den 27 november 1984 kl 10.15

Nils-Eric Sahlin

SECONDARY LEVELS IN DECISION MAKING

Lund 1984

This essay is a summary of the following six papers:

/1/ Preference among preferences as a method for obtaining a higher-ordered metric scale, The British Journal of Mathematical and Statistical Psychology, 34, 1981, pp. 62-75.

/2/ Unreliable probabilities, risk taking, and decision making, Synthese, 53, 1982, pp. 361-386. (With Peter Gärdenfors).

/3/ Decision making with unreliable probabilities, The British Journal of Mathematical and Statistical Psychology, 36, 1983, pp. 240-251. (With Peter Gärdenfors).

/4/ On second order probabilities and the notion of epistemic risk, in Foundations of Utility and Risk Theory with Applications, B.P. Stigum and F. Wenstøp (eds.), D. Reidel, 1983, pp. 95-104.

/5/ The role of second-order probabilities in decision making, in Analysing and Aiding Decision Processes, Humphreys, P.C., Svenson, O., and Vari, A. (eds.), Amsterdam, North-Holland, 1983, pp. 455-467. (With Robert Goldsmith).

/6/ Level of aspiration and risk, Philosophical Studies, No. 24, Department of Philosophy, Lund University, 1984.

I have had the opportunity to try out the ideas in this thesis in seminars and in papers read in a number of places, but also in discussions with friends and colleagues, and over the years many more people have helped me than I can possibly thank here. However, there are three people to whom I am grateful beyond measure. Peter Gärdenfors has been my intellectual guide from the very first course I took in philosophy. We have talked endlessly about the various topics of this thesis and it would probably never have been written had it not been for his unrestricted support. Sören Halldén introduced me to some of the subjects of the thesis and throughout my studies has encouraged, supported and offered me philosophical guidance. Robert Gold-

smith has taught me most of what I know about experimental psychological research, and has unselfishly assisted me in my own research.

Thanks to fellowships from the American-Scandinavian Foundation and the Swedish Academy of Sciences, I have been able to visit Princeton University, during the academic year 1979-1980, and Decision Research, Eugene, Oregon, during the spring of 1983. Excellent working conditions and stimulating discussions with friends and colleagues made it possible for me to outline and rewrite parts of this thesis during my stay at Princeton and Eugene.

I would also like to thank Bengt Hansson and Mari-Ann Saarnak for helping me with several practical matters, and Muriel Larsson for checking my English.

I dedicate this thesis to my wife Vivianne.

0. Introduction

Vivianne and I are planning to go to Copenhagen tomorrow. However, we have not yet decided whether we want to go. Our decision essentially depends on one important factor, the weather. Will it rain tomorrow? Since we want to spend our day in Copenhagen outdoors, we would stay at home, if we knew that it was going to rain. But it is not clear that it will be a showery day. If we decide to go and it starts to rain, the day will obviously be spoiled, but if the sun comes out, we will enjoy our day trip. If, on the other hand, we stay at home and it starts to rain, we can spend the day doing some useful but dull things, for example, cleaning the house. If it does not rain, we will certainly regret our decision.

What I have sketched above is a rather commonplace decision problem. It is a standard example of a type of problem which we face numerous times every day. Sometimes we stop and think about our possibilities, but most frequently we act without even considering our choice as part of a decision problem. However, for a decision theorist these types of elementary everyday problems are of special interest. In all their simplicity, they reveal the important aspects of decision making, aspects that any formal theory of decision making has to take into account and which will also be discussed in this summary.

In this summary I will go into some of the fundamental problems of decision making. The discussion will basically take place within the realm of what can be said to be a standard Bayesian theory of decision making, for example Savage's (1972) theory. Section 1 is therefore devoted to a presentation of the basic elements of this type of theory.



1. A Bayesian theory of probability and decision making

As I see it, decision making can be divided into four basic elements. That is, if you want to make a decision with the aid of a Bayesian decision theory (or any other decision theory for that matter), you have to structure the decision problem, assess probabilities, assess utilities, and calculate expected utilities. Each one of these four aspects of decision making deserves a somewhat fuller presentation.¹

The structure of a decision problem

A decision is a choice of one action alternative from a set of available alternatives, in a given situation. Let us, for simplicity, assume that this set of alternatives is finite and denote them $a_1, a_2, \dots, a_n$. For example, you might have a choice between buying a lottery ticket or not, or between going to Copenhagen or staying at home.

It is further assumed that the decision maker does not fully control the factors which determine the outcome of the decision. One can express this uncertainty by referring to different states of nature. For example, you do not yourself decide whether the lottery ticket you bought is a blank or whether it is going to rain in Copenhagen tomorrow.

If the decision maker has to take into consideration all possible states of nature, all possible developments of the world, we would have to assume an uncountable set of possible states. Therefore, for simplicity, we assume that in most decision situations of interest this uncountable set can be partitioned into a finite number of states relevant to the decision, denoted $s_1, s_2, \dots, s_m$.

Finally, the consequences or outcomes of choosing an action alternative a_i when the true state of nature turns out to be s_j is denoted o_{ij} . If you buy a lottery ticket and it turns out to be a blank, the outcome of your choice is that you have lost a sum of money. How much money you have lost depends on how much you paid for the lottery ticket. If we decide to go to Copenhagen and the weather is fine, then the consequence of our action is a pleasant day in Copenhagen.

Thus, if we want to solve a decision problem we have to determine

which states of nature, acts and outcomes, we want to consider. However, structuring a decision problem in this way is, of course, not enough in order to reach a decision, but it is the first step, and maybe the most important step, in the analysis of a specific problem.

Probabilities

For the decision maker it is important to know, not only that certain states are possible, but also exactly how likely it is that they will obtain. A Bayesian or subjectivist claims that these probabilities mirror the beliefs of the decision maker. Within strict Bayesianism it is assumed that these beliefs can be represented by a single probability measure defined over the states of nature. This assumption is very strong and its origin is to be found in the behavioural definition advocated by the Bayesian tradition.²

A Bayesian assumes that the decision maker's subjective probabilities are determined by his disposition to accept bets concerning certain states. It can be shown that the agent's degrees of belief satisfy the axioms of the probability calculus if one cannot construct a bet such that the agent will lose wealth independently of which state turns out to be the true one. That is, if such a bet cannot be found, then a unique probability measure exists which mirrors the agent's degrees of belief. Ramsey expresses this fact as follows:

Having degrees of belief obeying the laws of probability implies a further measure of consistency, namely such a consistency between the odds acceptable on different propositions as shall prevent a book being made against you.³

Ramsey's condition of consistency is more often referred to as the condition of coherence: Act in such a way that you 'prevent a book being made against you'.

In section 2 I will question the Bayesian assumption of the existence of a unique probability measure and thus, indirectly, the above behavioural definition of subjective probability. This does not mean that I reject the notion of subjective probability, far from it. As a basis for action subjective probability is the best tool we have.

Utilities

For the decision maker it is also important to know exactly how desirable the possible outcomes of his actions are. If a_i is chosen and s_j turns out to be the true state, what is the value of the corresponding outcome o_{ij} ?

Daniel Bernoulli (1730) forcefully argued that it is far from acceptable to evaluate the outcomes of our actions in terms of their monetary value. A lottery ticket, for example, clearly has a different value depending on whether you are 'rich' or 'poor'. The monetary value of a lottery ticket is the same whatever one's present level of wealth; however, its subjective value will differ from level to level. The problem that the monetary value of an outcome does not necessarily mirror its subjective value can be overcome by the introduction of a measure of utility.

A utility measure shall reflect, not only the ordinal preferences between the possible outcomes, but also the corresponding cardinal differences. It is not only of interest to know that I prefer coffee to tea after dinner, and tea to water, but also that I strongly prefer coffee to tea, but am rather indifferent to whether I get tea or water. The strength of my preferences is important for the decision I intend to make.

An interesting and fundamental question within the realm of decision making is how one can obtain knowledge about an agent's utility assessments. Atmospheric pressure can be measured by a barometer and heat by a thermometer. However, is there any reliable tool by which we can measure utilities? In 'Truth and probability' Ramsey presents one operational method for measuring utility. The method is based on comparisons among games. In section 3 I will present a new type of method for obtaining information about an agent's utility assessments. The method is based on the concept of preference among preferences. In section 4 both these methods will be employed in order to define a utility theory measure of risk.

Subjective expected utility

Assessing probabilities and utilities is not enough if one wants to reach a decision. The assessed values have to be combined in one way

or another if they are to be of any assistance to the decision maker.

The decision rule usually advocated (also by non-Bayesians) is based on the calculation of expected utility (or value). Consider, for example, a game which gives \$10 with the probability 0.50 and nothing otherwise. How much should we be prepared to pay in order to get the possibility of playing this game? What is a fair price? If we assume that dollars and utilities are exchangeable, the expected value of this game is \$5 ($=0.50 \times \$10 + 0.50 \times \0). Five dollars is a fair price. The expected value or utility of a game, lottery or act is obtained as the sum of the values or utilities of the outcomes that may result from each possible state of nature, multiplied by the corresponding state's probability:

$$u(a_i) = P(s_1)u(o_{i1}) + \dots + P(s_n)u(o_{in}).$$

In a decision problem there will be a number of acts to choose among. For example, there may be two games or lottery tickets and we have to select one of them. Consequently, each act will be associated with an expected utility value. The Bayesian theory of decision making defends the following decision rule:

Act in such a way that you maximize the subjective expected utility.

If one act has a higher expected utility value than another act, then the act which is associated with the higher value ought to be chosen. Comparing the game above with another game which gives \$20 with the probability 0.60 and -\$5 with the probability 0.40 (the minus sign indicates a loss of five dollars), we find that this latter game is preferable. The expected value of this game is \$10 which is greater than the \$5 value we computed above.

In section 2 I will advocate an alternative decision rule. The rule to be presented, however, has a great deal in common with the traditional Bayesian decision principle.

An alternative Bayesian decision theory

Savage's decision theory has been developed and reformulated in many

different ways. Others have also tried to find theories which can be thought of as serious alternatives to the decision model sketched above. One important Bayesian alternative is the theory suggested by Jeffrey (1965). A major difference between the traditional Bayesian decision theory and Jeffrey's theory is that the probability that a certain state of nature will occur is, in the latter theory, dependent on which act one chooses. Acting in one way or another may change the probability that certain outcomes will be obtained. Recent research has shown that there are greater differences between the two theories than was first thought and this fact has resulted in the development of a new type of decision theories, the so-called causal decision theories. However, there is another aspect of Jeffrey's theory which is of interest in this summary.

The so-called Bolker-Jeffrey axiomatization provides us with a system in which the agent's state of belief is represented by an infinite set of probability measures, rather than a unique measure. The result occurs because in this system the belief states are not determined by complete preference rankings of the underlying algebra; rather it is assumed that the agent can only provide a ranking of different subsets of the algebra. The agent may argue that coffee is preferred to tea and that wine is preferred to water, but we do not know, for example, if the agent prefers coffee to wine, or if water is preferred to coffee. Jeffrey's theory thus leads to the fact that the utility function is determined only up to a fractional linear transformation ($u'(\cdot) = (au(\cdot) + b) / (cu(\cdot) + d)$, for some conditions on the constants).⁴ In section 2 I will argue that Jeffrey's decision theory and the traditional Bayesian decision theory have the same problems, despite the fact that in the former theory states of beliefs are represented by sets of probability functions.

2. Second order probabilities and decision making

Let us assume that you are offered two games or lotteries and your task is to choose the one you consider to be most preferable.

Game A: In front of you on the table there is an opaque urn containing 30 black and 70 white balls. You are allowed to inspect the balls and the urn. In doing so you will find that the balls are identical except for their colour. If you draw a black ball from this urn, you receive \$100; however, if the drawing results in a white ball you will get nothing.

As the rational agent you are, I know that you will argue that the probability that you will draw a black ball is 0.30, $P(B)=0.30$, and, consequently, that the probability of a white ball is 0.70, $P(W)=0.70$. The expected value of this game is thus equal to \$30. If, for simplicity, we also assume that for small sums of money, dollars and utilities are exchangeable, the subjective expected utility of the game is 30 (utils). Thus, if you had to pay a sum of money in order to be allowed to play this game, it could be argued that \$30 would be a fair price.

The second game we shall consider is somewhat different.

Game B: A classic event is the annual Oxford versus Cambridge boat race. Let us say that you receive \$100 if Cambridge wins and nothing if Oxford wins.

We assume that from knowledge of the two teams' prior results, you believe that the probability that Cambridge will win is 0.30, $P(C)=0.30$, and, thus, that the probability that Oxford will gain the victory is 0.70, $P(O)=0.70$. This implies that the expected value of this second game equals the expected value of the first game, or that the subjective expected utility of the two games is equal. Also in this case 30 dollars would be a fair price for the game.⁵

However, although the two games have the same value for you, I am pretty sure that faced with the task of choosing one of them, you will take game A. Or, for the sake of the argument let us assume that you make this choice. But, isn't this a choice contrary to the strict

Bayesian decision theory, a decision contrary to the principle of maximizing subjective expected utility (value)? The two games have the same subjective expected utility. Therefore, B is as preferable as A.

Why do you prefer A? Is it because your assessment of the probability of C, that Cambridge will win, is too high? Let us assume it is not; $P(C)=0.30$, 0.30 is the best assessment you can make. The crux is that you feel that you know more about the first game situation than about the second. There is a difference in degree between your ignorance concerning drawings from urns and boat races. You feel that you know much more about the probability distribution in game A than in game B. In the second game we are dealing with what could be called indeterminate or unreliable probabilities. In game A this unreliability can be ignored.

Actually, I believe that there is no way by which you can obtain information about the boat race, which will make you feel as easy about the probability assessment in game B, as you are about the assessment that a black ball will be drawn. One problem about the strict Bayesian or subjectivistic view of probability is that it cannot account for this obvious difference in ignorance which prevails in the two game situations. Savage is certainly clear about the problem:

To approach the matter in a somewhat different way, there seem to be some probability relations about which we feel relatively 'sure' as compared with others. When our opinions, as reflected in real or envisaged action, are inconsistent, we sacrifice the unsure opinions to the sure ones. The notion of 'sure' and 'unsure' introduced here is vague, and my complaint is precisely that neither the theory of personal probability, as it is developed in this book, nor any other device known to me renders the notion less vague.⁶

The strict Bayesian or subjectivistic approach to probability also has another problem, closely related to the one discussed above. It is mentioned by Savage in the following way:

A second difficulty, perhaps closely associated with the first one, stems from the vagueness associated with judgments of the magnitude of personal probability. The postulates of

personal probability imply that I can determine, to any degree of accuracy whatsoever, the probability (for me) that the next president will be a Democrat.⁷

What Savage alludes to is the fact that your (and my) belief in the statement 'the next president (of the United States) will be a Democrat' can be fixed to any degree we like. This is a central theorem of the theory. But, isn't this asking too much? Consider, for example, an urn containing one million balls of which 540 720 are black, the remainder white. You should, thus, argue that the probability that you will draw a black ball from the urn is 0.540720, an assessment which seems fine for this type of unrealistic situation. However, are there any real-life events for which you can assess the probability with the same accuracy? What is the probability of rain in Copenhagen tomorrow? What is the probability that the next president of the United States will be a Democrat? What is the probability that there will be a transit strike in Verona, Italy next week?

My guess is that you will answer each one of these questions by a one place or two place decimal number. I can probably squeeze a third and a fourth decimal number out of you, but you will emphasise that these assessments are very unreliable. A fifth and a sixth decimal assessment is out of the question.

These examples show that the ordinary Bayesian assumption that the agent's state of belief concerning which state obtains can be represented by a unique probability measure, and to any degree of precision, is too strong and ought to be relaxed.⁸ I will now suggest one way by which the above difficulties can be solved, a solution which results in the formulation of a new type of decision theory.

A decision theory

The theory presented in /2/ and discussed and further developed in /3/, /4/, /5/ and Sahlin (1984b), assumes that the agent's knowledge and beliefs about the relevant states can be represented by a class P of probability measures, the set of all epistemically possible probability measures, i.e. those and only those measures which do not contradict the knowledge of the agent. According to the theory, each state of nature is associated with a set of possible values $P(s_i)$,

where $P(.)$ belongs to P .

However, using all epistemically possible measures as a basis for action is hardly realistic. Observing that some probability measures are more reliable than others will make us discriminate between them. Perhaps you believe that the probability that it will rain tomorrow is close to 0.20, but, for all you know, it is epistemically possible that the probability is 0.10, 0.50 or 0.90. However, you feel that a 0.20 assessment is much more reliable than a 0.10 or a 0.50 assessment, which in turn are more reliable than a 0.90 assessment. Or, in Savage's words, 0.20 is a sure option or probability compared with 0.10 or 0.50.

One can formally model this feeling of reliability concerning one's probability assessments by introducing a second order probability measure $r(.)$, a measure of epistemic reliability, defined over the set P of epistemically possible probability measures. The measure is intended to reflect how complete or adequate the knowledge is, upon which one's first order probability assessment is based.

The theory has an ordered set of rules for reaching a decision. Firstly, the set P is restricted to a set P/r_0 of probability measures for which the level of epistemic reliability is deemed satisfactory, i.e. a subset P/r_0 of the set P is selected such that for all $P(.)$ belonging to P/r_0 , $r(P(.)) \geq r_0$. Secondly, the expected utility of each alternative a_i and of each distribution $P(.)$ in P/r_0 is computed and the minimal expected utility of each alternative determined. Finally, the action alternative with the largest minimal expected utility is selected.⁹

It is easily seen that a theory like the one sketched here avoids the type of problems discussed above. The agent is not required to act upon a unique probability distribution. For example, we might feel that it is highly reliable that the probability that the next president of the United States will be a Democrat is between 0.40 and 0.50. This means that our set P/r_0 of epistemically possible and satisfactorily reliable probability distributions consists of those and only those distributions for which $0.40 \leq P(\text{the next president of the United States will be a Democrat}) \leq 0.50$. Thus, in a straightforward way, what Savage calls 'the vagueness associated with judgments of the magnitude of personal probability' is captured by the model.

The first problem we discussed, the fact that under some conditions we seem to 'sacrifice the unsure opinions to the sure ones', we prefer, for example, situations where the assessed probability distribution is deemed as reliable, is solved by the second order probability measure introduced by the theory. In order to see this let us once again consider game B. If the agent is given the chance of using more than one probability distribution, this will most likely result in assessment of a set P/r_0 of satisfactorily reliable measures.

Let us, for simplicity, assume that this set can be represented as the probability interval (0.20, 0.40). Now applying the present decision theory we find that the maximal minimal expected utility of game B is \$20, which is clearly less than \$30. This explains why game A is considered to be preferable to game B, despite the fact that, according to traditional Bayesian decision theory, the two games have the same expected utility. The fact of the matter is that unrealistic demands made on the agent force him or her to make a seemingly irrational choice, a choice which, in view of the present theory, is perfectly rational and reasonable.

In order to emphasize why a theory based on higher order probabilities is preferable to the traditional type of theories let us consider a very much discussed example or problem. Our urn in this case is supposed to contain exactly fifty percent white balls and fifty percent black balls, meaning that our state of knowledge is such that we believe that the probability of drawing a black ball is a half. We are, however, allowed to draw one ball at a time from the urn, inspect the colour of the ball, and then return it. Let us say that we have drawn one million balls and that half of them are white. According to strict Bayesian theory, the conditional probability that the next ball will be a white ball is still a half. Our degree of belief is thus completely insensitive to the vast amount of evidence we have received. This problem is known as the problem of total evidence or Popper's paradox of ideal evidence.

It is, however, easily seen that with the aid of the present model of an agent's state of belief there is no paradox. Although we do not change our first order probability assessments, the probability that we will draw a black ball is still a half, we do change our second order probability assessments, i.e. we modify the epistemic reliabil-

ity assessment. When one million balls are drawn from the urn, of which fifty percent are black, we are very 'sure' that the probability that the next ball will be black is really a half, $r(P(B)=0.50)$ is close to maximal.

With the aid of the present model of an agent's state of belief a new and, I believe, important notion of risk can be introduced and defined. If in game B above you decide only to consider the probability value 0.30 ($P(\text{Cambridge win})=0.30$), you are taking a risk, an epistemic risk. The reason why you feel that you are taking a risk is because you feel that a 0.30 assessment is fairly unreliable. There are many other values which are almost as reliable and thus maybe should be considered. Choosing a high r_0 -level means that the restriction of P to P/r_0 is considerable, a high epistemic risk is taken, many epistemically possible first order probability distributions are not considered when taking the decision. On the other hand, if a low r_0 -level is chosen this generally implies that a lower epistemic risk is taken. The notion of epistemic risk is discussed in /4/, but also in /2/ and /5/.

Recently a number of authors have tried to develop theories which avoid the problems discussed above. One prominent theory is the one developed by Levi (1974, 1980). Another interesting theory has been suggested by Kyburg (1983b). The two theories attempt to replace the assessment of a unique probability measure by an assessment of a set of such measures (Levi argues that this set has to be convex. Kyburg started out by interval assessments, but has later revised his views concerning such assessments and has replaced them by sets of probability measures). There are several significant differences between the theory presented here and these two theories (see /2/, Gärdenfors & Sahlin (1982b), Levi (1982) and Sahlin (1984b)). Each of the three theories advocates a decision rule distinct from the rule of the other theories. Arguments in favour of the principle presented above can be found in /2/ and Sahlin (1984b). However, the major difference, I believe, is that the present theory introduces a second order probability measure, the measure of epistemic reliability.

In a note in the second revised edition of The Foundations of Statistics, Savage mentions the possibility of using sets of probability measures rather than a single measure:

One tempting representation of the unsure is to replace the person's single probability measure P by a set of such measures, especially a convex set.¹⁰

However, Savage is clearly against higher order probabilities. In /4/ Savage's arguments and some other standard arguments against higher order probabilities are discussed and rejected.

In /3/ the present decision theory is further developed and evaluated in the light of empirical evidence on ambiguity and uncertainty in decision making. Paper /5/ presents two descriptive models of the use of second order probabilities. Experimental studies supported the assumption that a majority of subjects, in one way or another, make use of second order probabilities in their decision making. Second order probability is thus an important concept for descriptive as well as normative theories.

It is now also obvious why Jeffrey's theory fares no better than Savage's theory with respect to the problem discussed above.¹¹ The set of probability measures produced by Jeffrey's theory is the result of the agent's inability to express fully strength of preference. However, the theory does not take the indeterminacy into account. The clarity of perception of uncertainty is not introduced by the theory. We see this since this clarity would correspond in Jeffrey's theory to a clarity of perception of preference, not to an inability to express fully strength of preference. Thus, Jeffrey's theory has problems very similar to the traditional Bayesian theory, which is hardly surprising since Jeffrey's theory can be viewed as a modified Bayesian theory.

3. Second order preferences

It is commonly agreed that our preferences reveal our utility assignments. In this section, I intend to discuss two methods by which we can obtain information about an individual's utility assignments through a study of his or her preferences.

Let us assume that o_1, o_2, o_3 and o_4 are four possible outcomes such that the agent prefers o_1 to o_2 , o_2 to o_3 and o_3 to o_4 . These preferences can be modelled either by a weak preference relation 'it is better than or equal to' or by a strict relation 'it is strictly better than'. In /1/ I employed the former type of relation. In this summary, however, I take B to denote the strict binary preference relation.

We can then more formally write the above stated preferences as: o_1Bo_2, o_2Bo_3 and o_3Bo_4 . Assuming that preferences mirror utilities we know that o_1 has to be assigned a utility value higher than that assigned to o_2 , that o_1Bo_2 indicates that $u(o_1) > u(o_2)$, and similarly for o_2 and o_3 , and for o_3 and o_4 . We thus know that $u(o_1) > u(o_2) > u(o_3) > u(o_4)$.

But the above preferences do not reveal any cardinal information whatsoever. We do not know if $u(o_1)-u(o_2) \geq u(o_3)-u(o_4)$, or if, for example, the reverse inequality holds. We have not obtained any information about the value distances between the four objects. Such information is, of course, necessary if the utility values are to be used in more complex decision situations. Ramsey (1978) suggests the following method by which the problem can be solved.

Let us assume that a proposition or event E exists such that $P(E)=P(-E)=1/2$. The agent believes that the probability that E will occur equals the probability that E will not occur. In Ramsey's terminology E should be an 'ethically neutral proposition'. It might prove difficult to find an event E which in practice satisfies the necessary conditions; however, let us assume that a fair coin, or something like that, can be employed to simulate an ethically neutral proposition.

We now assume that the agent is offered one of the following two fifty-fifty games:

Game A: If E occurs o_1 is received.
 If $\neg E$ occurs o_4 is received.

Game B: If E occurs o_2 is received.
 If $\neg E$ occurs o_3 is received.

If the agent prefers game A to game B we know that $u(A)$ is greater than $u(B)$. This fact implies that $P(E)u(o_1) + P(\neg E)u(o_4) > P(E)u(o_2) + P(\neg E)u(o_3)$ and since $P(E) = P(\neg E)$ it follows that $u(o_1) - u(o_2) > u(o_3) - u(o_4)$. The preference for game A reveals that the value distance between o_1 and o_2 is greater than the corresponding distance between o_3 and o_4 . The information obtained can be illustrated heuristically as follows:

Best o_1 ----- o_2 o_3 ----- o_4 Worst

That there is no line connecting o_2 with o_3 simply indicates that so far we do not know the relative length of this value distance. If we want to find out, we have to employ two other games.

If Ramsey's method is employed throughout the entire set of outcomes a complete ordering can be obtained. Furthermore, as Ramsey shows, under certain conditions one can determine the utility function up to a positive linear transformation by using the above method. We can summarize Ramsey's method as follows:

Def R: If o_1 is better than o_2 and o_3 is better than o_4 , then the value distance between o_1 and o_2 is greater than the value distance between o_3 and o_4 , if and only if, it is better to play a fifty-fifty game between o_1 and o_4 than to play a fifty-fifty game between o_2 and o_3 .

Ramsey's method has been questioned for different reasons. The core of the criticism derives its force from the fact that the method is based on comparisons among games. The agent's attitude towards gambling will most likely penetrate his or her preferences. It may be the case that we do not reveal the agent's true value distances, but rather value distances affected by his or her attitude towards risk.

In order to answer this criticism, Ramsey's method has to be compared with an alternative method not making use of games or lotteries. In /1/ such a method is presented and discussed in relation to Ramsey's method.¹²

The new method is based on the concept of second order preferences or preference among preferences, and it was first introduced by Halldén (1980).

In addition to the normative preference relation B above, we assume the existence of a descriptive preference relation P. The sentence o_1Po_2 compares the entities o_1 and o_2 and asserts that if it is possible for the agent to realize o_1 or o_2 , o_1 will be realized. The measuring procedure discussed in /1/ can now be formulated as follows.

Let us assume that the ordinal ordering between the four outcomes o_1 , o_2 , o_3 and o_4 is the same as above. However, instead of a choice between two games (A and B) the agent is asked the following question:

Is it better to be able to prefer
 o_1 to o_2
 than to prefer
 o_3 to o_4 ?

Put in terms of the secondary preference relation the agent is asked whether $(o_1Po_2)B(o_3Po_4)$ or if $(o_3Po_4)B(o_1Po_2)$.

If the former is the case, if it is better to prefer o_1 to o_2 , we say that the value distance between o_1 and o_2 is greater than the value distance between o_3 and o_4 , $u(o_1)-u(o_2) > u(o_3)-u(o_4)$.

The method is summarized in the following definition:

Def H: If o_1 is better than o_2 and o_3 is better than o_4 , then the value distance between o_1 and o_2 is greater than the value distance between o_3 and o_4 , if and only if, it is better to prefer o_1 to o_2 than to prefer o_3 to o_4 .

The idea behind this definition is that strength of preference is

revealed by one's second order preferences. If my mental powers for one reason or another were to be restricted, for example due to a brain operation or simply because of a growing state of senility, and I am given the possibility to select a number of preferences that I have to give up, then it is important that I make the selection in such a way that I will still be able to state serious preferences -- greater value distances are important. How this selection should be made is reflected through my second order preferences.

Some illuminating examples of how the definition works in practice can be found in /1/ and Sahlin (1982); hence I refrain from discussing them here. However, the results of /1/ should be mentioned.

In /1/ it is shown how def H can be used in practice and it is shown that a higher ordered metric scale can be obtained by employing this definition. The latter result means that we can obtain very precise information concerning the agent's value distances. The method presented involves a minimum of comparisons. It is also shown that if def H and def R are employed to obtain orderings of value distances the discrepancies between the two orderings are few. It is argued that, for each subject, the differences between the two orderings thus obtained can be attributed to two main causes: risk aversion and risk proneness, both in connexion with the use of def R.

Writing /1/, however, I did not fully realize the importance of these discrepancies between the value distances obtained. My present view is that they are of interest, for example, for the development of a theory of risk, a problem which we will now turn to.

4. Risk

The seventies were the decade when we suddenly realized that we live in a hazardous world; there are a lot of risks to be accounted for. We have the risk of a nuclear war, the risk of a nuclear power plant accident, the risk of recombinant DNA research, and the risk of losing two Sw.Cr. on the football pools. However, the concept of risk proved to be extremely difficult to analyse and, consequently, a vast number of rather independent theories of risk appeared on the market. In /6/ my aim is to analyse one type of theory of risk, the one developed within the framework of modern utility theory.

In two important works by Arrow (1971) and Pratt (1964), it is suggested that risk can be defined in terms of the agent's utility function. Assume for example, that you are offered either \$50 for sure or a lottery ticket which gives \$100 with the probability 0.50 and nothing otherwise, and that you prefer the fifty dollars. According to Arrow and Pratt, this type of behaviour is a sign of risk aversion. The preference also indicates (see section 3) that the value distance between \$100 and \$50 is less than the value distance between \$50 and \$0.

Extending this basic idea to the continuous case, we find that an agent is risk averse if and only if the utility function is concave. How risk averse the agent is will depend on the concavity of his or her utility function. Thus, we simply need to measure the curvature of the agent's utility function in order to know how risk averse or risk prone he or she is. Arrow and Pratt therefore suggest the following measure as a possible measure of risk:

$$r(x) = -u''(x)/u'(x),$$

where $u(\cdot)$ is assumed to be twice continuously differentiable.

However, from a number of different standpoints it can be argued that the measure $r(\cdot)$ is not an acceptable measure of risk.¹³ In /6/ I try to extend these arguments against $r(\cdot)$ and show that it is not an adequate measure of risk and thus has to be revised.

Following Arrow and Pratt, an agent is, for example, risk prone as soon as he or she prefers a fair game to its cash equivalent. However, not all such preferences may be a sign of risk proneness. Assume that

you need \$50. If I offer you either \$100 for sure or a lottery ticket which gives \$150 with the probability 0.50 and \$50 with the probability 0.50, and you prefer the lottery ticket to \$100 for sure (which is the cash equivalent of the game), then you are, according to Arrow's and Pratt's theory, risk prone. But whatever you choose, you will satisfy your need. If the goal for which you strive can be identified with a level of aspiration, it may be said that you will satisfy your level of aspiration whatever you do. The example shows that risk must be measured in situations which can be thought of as involving factual risk. In /6/ the following definition of factual risk is suggested:

An action is said to involve factual risk if the action has a set of possible specific outcomes and for this outcome set at least one but not all outcomes are strictly below the level of aspiration of the decision situation.

Arrow's and Pratt's measure of risk can now be revised as follows. Let $u_0(.)$ be a utility function obtained by a method like def H, a method which does not involve comparisons among games. One could say that this function is the agent's true or real utility function. Similarly, we let $u_1(.)$ be a utility function obtained by use of a method like def R, but with the restriction made to factual risk situations. This utility function can thus be assumed to be affected by the agent's attitude towards risk. It is, however, argued that none of these functions on their own can tell us anything about risk, but comparing the two utility functions will give us the information we are looking for.

If $u_1(.)$ is different from $u_0(.)$ (disregarding positive linear transformations) a shift in the agent's value distances has taken place, a shift which can be attributed to the risk perception of the agent. If we let $r_1(x) = -u_1'(x)/u_1'(x)$ and $r_0(x) = -u_0''(x)/u_0'(x)$, then we can capture this idea by defining risk as follows:

$$R(x) = r_1(x) - r_0(x),$$

where $u_1'(x)$, $u_0'(x) \neq 0$. If $u_1'(x)$ or $u_0'(x) = 0$ we let $r_1(x)$ and $r_0(x) = 0$, respectively.

Compared with Arrow's and Pratt's measure of risk it, can be shown that $R(\cdot)$ avoids some of the criticism levelled against this measure. For example, a choice of the lottery ticket in the above example is not necessarily a sign of risk proneness. If the utility function $u_o(\cdot)$ is such that the value distance between \$50 and \$100 is less than the value distance between \$100 and \$150, then the choice is perfectly rational and does not reveal any risk proneness. The agent's choice is in accordance with his or her true value distances. However, if the agent considers these two value distances to be equal, or that the former is greater than the latter, $u_o(\$150) - u_o(\$100) \leq u_o(\$100) - u_o(\$50)$, then a shift in his or her value distances has obviously taken place since $u_i(\$150) - u_i(\$100) > u_i(\$100) - u_i(\$50)$. This shift is most adequately explained by referring to the agent's risk proneness, a fact which is captured by $R(\cdot)$, but not by $r(\cdot)$.

One could describe the major difference between the present theory and Arrow's and Pratt's theory by saying that for the latter theory, risk is measured relative to monetary outcomes, whereas in the present theory, risk is measured relative to the agent's utility function.

Notes

- 1) Completely developed decision theories can be found in, for example, Halldén (1980), Jeffrey (1965) and Savage (1972).
- 2) See Ramsey (1978), but also Savage (1972). For criticism of this assumption, see Kyburg (1983a) and Levi (1974, 1980).
- 3) See Ramsey (1978), p. 85.
- 4) See Jeffrey (1965, 1983).
- 5) If you disagree with me concerning the probability assessment in Game B, I will simply change the composition of the urn in such a way that the expected utility of the two games is equal.
- 6) Savage (1972), pp. 57-58.
- 7) Savage (1972), p. 59.
- 8) Kyburg (1983a) has discussed these problems. However, he seems to have a somewhat different intuition than I have. See also Levi (1974, 1980).
- 9) This final decision rule was first suggested by Gärdenfors (1979).
- 10) Savage (1972), p. 58, note +.
- 11) Jeffrey is obviously arguing that it does in his (1983).
- 12) Ramsey's method and its problems are discussed in Sahlin (1982).
- 13) See Hansson (1977, 1981).

References

- Arrow, K.J. Essays in the Theory of Risk-Bearing. North Holland, London, 1971.
- Bernoulli, D. Exposition of a new theory of measurement. Econometrica, 1954 (1730), 22, 23-36.
- Goldsmith, R.W., & Sahlin, N.-E. The role of second-order probabilities in decision making. In Humphreys, P.C., Svenson, O., & Vari, A. (eds.), Analysing and Aiding Decision Processes, North-Holland, Amsterdam, 1983, 455-467.
- Gärdenfors, P. Forecasts, decisions and unreliable probabilities. Erkenntnis, 1979, 14, 159-181.
- Gärdenfors, P., & Sahlin, N.-E. Unreliable probabilities, risk taking and decision making. Synthese, 1982a, 53, 361-386.
- Gärdenfors, P., & Sahlin, N.-E. Reply to Levi. Synthese, 1982b, 53, 433-438.
- Gärdenfors, P., & Sahlin, N.-E. Decision making with unreliable probabilities. The British Journal of Mathematical and Statistical Psychology, 1983, 36, 240-251.
- Halldén, S. The Foundations of Decision Logic. Library of Theoria, Lund, 1980.
- Hansson, B. Risk aversion as a problem of conjoint measurement. Unpublished manuscript, Lund, 1980.
- Hansson, B. The decision game -- The conceptualisation of risk and utility. In Ethics, Proceedings of the 5th international Wittgenstein symposium, 1981, 187-193.
- Jeffrey, R.C. The Logic of Decision, McGraw-Hill, New York, 1965.

- Jeffrey, R.C. Bayesianism with a human face. In Earman, J. (ed.), Testing Scientific Theories, Minnesota Studies in the Philosophy of Science, X, University of Minnesota Press, Minneapolis, 1983, 133-156.
- Kyburg, H.E. Epistemology and Inference. University of Minnesota Press, Minneapolis, 1983a.
- Kyburg, H.E. Rational belief. The Behavioral and Brain Sciences, 6, 1983b, 231-273.
- Levi, I. On indeterminate probabilities. Journal of Philosophy, 1974, 71, 391-418.
- Levi, I. The Enterprise of Knowledge. MIT Press, Cambridge, Mass., 1980.
- Levi, I. Ignorance, probability, and rational choice. Synthese, 1982, 53, 387-417.
- Pratt, J.W. Risk aversion in the small and in the large. Econometrica, 1964, 32, 122-136.
- Ramsey, F.P. Truth and probability. In Foundations, (ed.), D.H. Mellor, Routledge & Kegan Paul, London, 1978.
- Sahlin, N.-E. Preference among preferences as a method for obtaining a higher-ordered metric scale. British Journal of Mathematical and Statistical Psychology, 1981, 34, 62-75.
- Sahlin, N.-E. Ramseys sannolikhetsfilosofi. Philosophical Studies, No. 16, Department of Philosophy, Lund University, 1982.
- Sahlin, N.-E. On second order probabilities and the notion of epistemic risk. In Stigum, B.P., & Wenstøp, F. (eds.), Foundations of Utility and Risk Theory with Applications, D. Reidel, 1983, 95-104.

Sahlin, N.-E. Level of aspiration and risk. Philosophical Studies, No. 24, Department of Philosophy, Lund University, 1984a.

Sahlin, N.-E. Three decision rules for generalized probability representation. Lund, 1984b.

Savage, L.J. The Foundation of Statistics. Dover, New York, 1972.

von Neumann, J., & Morgenstern, O. Theory of Games and Economic Behaviour. Princeton University Press, Princeton, 1944.

